CLAIMS

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is as follows:

1. A method of resource allocation to yield a benefit comprising the steps of:

associating each customer's demand with a benefit gained; and

finding a time-varying resource allocation that would yield a benefit .

2. The method of resource allocation as recited in claim 1, further comprising the steps of:

discounting future benefits; and

finding optimal allocations of resources from current time through current time plus lookahead based on discounted benefit and forecast demand, wherein the step of discounting future benefits is based on a future discounting algorithm.

3. The method of resource allocation recited in claim 2, wherein the future discounting algorithm is based on a policy which, when faced with a choice between a guaranteed benefit immediately and a potential benefit in the future, a decision is made by comparing the guaranteed benefit value with a discounted value of the potential future benefit.

4. The method of resource allocation recited in claim 2, wherein the future discounting algorithm is a deterministic algorithm that achieves a competitive ratio of $\left(1+\frac{1}{L}\right)\sqrt[L]{L+1}$, where L is a lookahead factor which models some amount of future demand known to a provider of the resource.

5. The method of resource allocation recited in claim 2, wherein the algorithm is an intermittent reset algorithm that achieves a competitive ratio of $1+\frac{4}{(L-7)}$, where L is a lookahead factor which models some amount of future demand known to a provider of the resource.

6. The method of resource allocation as recited in claim 1, wherein resource allocation is done to maximize a benefit.

7. The method of resource allocation as recited in claim 1, wherein the benefit is a tangible benefit.

8. The method of resource allocation as recited in claim 7, wherein the tangible benefit is a profit and resource allocation is done to maximize the profit.

9. The method of resource allocation as recited in claim 1, wherein the benefit is an intangible benefit.

10. The method of resource allocation as recited in claim 9, wherein the intangible benefit is customer satisfaction and resource allocation is done to maximize customer satisfaction.

11. The method of resource allocation as recited in claim 1, wherein the resource is computer cycles and resource allocation is done to more efficiently solve computationally intensive problems.

12. A method of resource allocation to yield a benefit comprising the steps of:

modeling the resource allocation problem mathematically;

in the model, dividing time into intervals of fixed length based on the assumption that demand is uniformly spread throughout each such interval; and

associating each customer's demand with a benefit gained and finding a time-varying resource allocation that would maximize benefit gained.

13. The method of resource allocation as recited in claim 12, further comprising the steps of:

discounting future benefits; and

finding optimal allocations of resources from current time through current time plus lookahead based on discounted benefit and forecast demand, wherein the step of discounting future benefits is based on a future discounting algorithm.

14. The method of resource allocation recited in claim 13, wherein the future discounting algorithm is based on a policy which, when faced with a choice between a guaranteed benefit immediately and a potential benefit in the future, a decision is made by comparing the guaranteed benefit value with a discounted value of the potential future benefit.

15. The method of resource allocation recited in claim 13, wherein the future discounting algorithm is a deterministic algorithm that achieves a competitive ratio of $\left(1+\frac{1}{L}\right)\sqrt[L]{L+1}$, where L is a lookahead factor which models some

amount of future demand known to a provider of the resource.

16. The method of resource allocation recited in claim 12, wherein the future discounting algorithm is an intermittent reset algorithm that achieves a competitive ratio of $1+\frac{4}{(L-7)}$, where L is a lookahead factor which models some amount of future demand known to a provider of the resource.

17. The method of resource allocation as recited in claim 12, wherein the benefit is a tangible benefit.

18. The method of resource allocation as recited in claim 17, wherein the tangible benefit is a profit and resource allocation is done to maximize the profit.

19. The method of resource allocation as recited in claim 12, wherein the benefit is an intangible benefit.

20. The method of resource allocation as recited in claim 19, wherein the intangible benefit is customer satisfaction and resource allocation is done to maximize customer satisfaction.

21. The method of resource allocation as recited in claim 12, wherein the resource is computer cycles and resource allocation is done to more efficiently solve computationally intensive problems.

22. A method for server allocation in a Web server "farm" based on limited information regarding future loads to achieve close to greatest possible

revenue based on an assumption that revenue is proportional to the utilization of servers and differentiated by customer class comprising the steps of:

modeling the server allocation problem mathematically;

in the model, dividing time into intervals of fixed length based on the assumption that each site's demand is uniformly spread throughout each such interval;

maintaining server allocations fixed for the duration of an interval, servers being reallocated only at the beginning of an interval, and a reallocated server being unavailable for the length of the interval during which it is reallocated providing time to "scrub" the old site (customer data) to which the server was allocated, to reboot the server and to load the new site to which the server has been allocated, each server having a rate of requests it can serve in a time interval and customers share servers only in the sense of using the same servers at different times, but do not use the same servers at the same time; and

associating each customer's demand with a benefit gained by the service provider in case a unit demand is satisfied and finding a time-varying server allocation that would maximize benefit gained by satisfying sites' demand.

23. The method for server allocation in a Web server "farm" as recited in claim 22, further comprising the steps of:

discounting future benefits; and

finding optimal allocations of servers from current time through current time plus lookahead based on discounted revenues and forecast demand, wherein the step of discounting future benefits is based on a future discounting algorithm.

24. The method for server allocation in a Web server "farm" as recited in claim 23, wherein the future discounting algorithm is based on a policy which, when faced with a choice between a guaranteed benefit immediately and a potential benefit in the future, a decision is made by comparing the guaranteed benefit value with a discounted value of the potential future benefit.

25. The method for server allocation in a Web server "farm" as recited in claim 22, wherein the future discounting algorithm is a deterministic algorithm that achieves a competitive ratio of $\left(1+\frac{1}{L}\right)\sqrt[L]{L+1}$, where L is a lookahead factor which models some amount of future demand known to a provider of the resource.

26. The method for server allocation in a Web server "farm" as recited in claim 22, wherein the future discounting algorithm is an intermittent reset algorithm that achieves a competitive ratio of $1+\frac{4}{(L-7)}$, where L is a lookahead factor which models some amount of future demand known to a provider of the resource.

27. The method for server allocation in a Web server "farm" as recited in claim 23, wherein resource allocation is done to maximize profit.